Object Detection Through Exploration With A Foveated Visual Field

Emre Akbas, Miguel P. Eckstein

Abstract—We present a foveated object detector (FOD) as a biologically-inspired alternative to the sliding window (SW) approach which is the dominant method of search in computer vision object detection. Similar to the human visual system, the FOD has higher resolution at the fovea and lower resolution at the visual periphery. Consequently, more computational resources are allocated at the fovea and relatively fewer at the periphery. The FOD processes the entire scene, uses retino-specific object detection classifiers to guide eye movements, aligns its fovea with regions of interest in the input image and integrates observations across multiple fixations. Our approach combines modern object detectors from computer vision with a recent model of peripheral pooling regions found at the V1 layer of the human visual system. We assessed various eye movement strategies on the PASCAL VOC 2007 dataset and show that the FOD performs on par with the SW detector while bringing significant computational cost savings.

Index Terms—object detection, visual search, eye movements, latent linear discriminant analysis, foveated visual field

arXiv:1408.0814v1 [cs.CV] 4 Aug 2014

1 Introduction

There has been substantial progress (e.g. [7], [14], [22], [29], [37], [46], [53] to name a few) in object detection research in recent years. However, humans are still unsurpassed in their ability to search for objects in visual scenes. The human brain relies on a variety of strategies [9] including prior probabilities of object occurrence, global scene statistics [33], [45] and object co-occurrence [10], [28], [35] to successfully detect objects in cluttered scenes. Object detection approaches have increasingly included some of the human strategies [1], [2], [11], [14], [41]. One remaining crucial difference between the human visual system and a modern object detector is that while humans process the visual field with decreasing resolution away [26], [39], [43], [49] from the fixation point and make saccades to collect information, typical object detectors [14] scan all locations at the same resolution and repeats this at multiple scales. The goal of the present work is to investigate the impact on object detector performance of using a foveated visual field and saccade exploration rather than the dominant sliding window paradigm [14], [29], [53]. Such endeavor is of interest for two reasons. . First, from the computer vision perspective, using a visual field with varying resolution might lead to reduction in computational complexity, consequently the approach might lead to more efficient object detection algorithms. Second, from a scientific perspective, if a foveated object detection model can achieve similar performance accuracy as a non-foveated sliding window approach, it might suggest a possible reason for the evolution of foveated systems in organisms: achieving successful object detection while minimizing computational and metabolic costs.

Contemporary object detection research can be roughly outlined by the following three important components of a modern object detector: the features, the detection model and the search model. The most popular choices for these three components are Histogram of Oriented Gradients (HOG) features [6], mixture of linear templates [14], and the sliding window (SW) method, respectively. Although there are efforts to go beyond these standard choices (e.g. new features [37], [46]; alternative detection models [22], [46], whether object parts should be modeled or not [8], [54]; and alternative search methods [13], [14], [20], [24], [46]), HOG, mixture of linear templates and SW form the crux of modern object detection methods ([14], [29], [53]). Here, we build upon the "HOG + mixture of linear templates" framework and propose a biologically inspired alternative search model to the sliding window method, where the detector searches for the object by making saccades instead of processing all locations at fine spatial resolution (See Section 4 for a more detailed discussion on related work).

The human visual system is known to have a varying resolution visual field. The fovea has higher resolution and this resolution decreases towards the periphery [26], [39], [43], [49]. As a consequence, the visual input at and around the fixation location has more detail relative to peripheral locations away from the fixation point. Humans and other mammals make saccades to align their high resolution fovea with the regions of interest in the visual environment. There are many possible methods to implement such a foveated visual field in an object detection system. In this work, we opt to use a recent model [15] which specifies how responses of elementary sensors are pooled at the layers (V1 and V2) of the human visual cortex. The model specifies the shapes and sizes of V1, V2 regions which pool responses from the visual field. We use a simplified version of this

• *E. Akbas and M. P. Eckstein are with the Vision and Image Understanding Laboratory at the Department of Psychological and Brain Sciences, University of California Santa Barbara.*

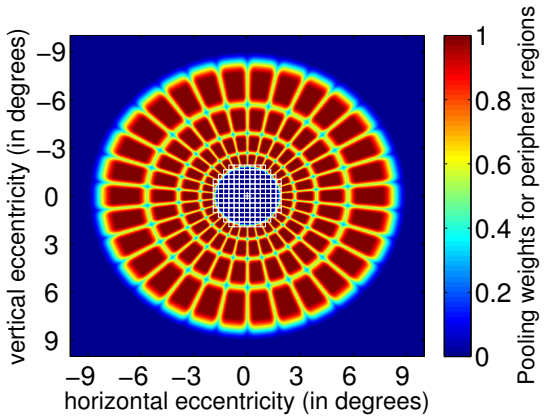

Fig. 1. The foveated visual field of the proposed object detector. Square blue boxes with white borders at the center are foveal pooling regions. Around them are peripheral pooling regions which are radially elongated. The sizes of peripheral regions increase with distance to the fixation point which is at the center of the fovea. The color within the peripheral regions represent pooling weights.

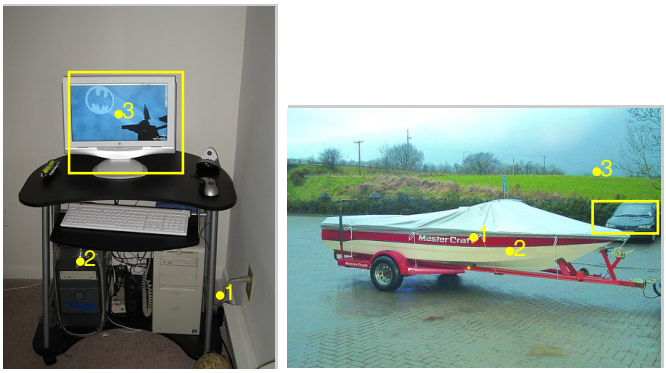

Fig. 2. Two example detections by our foveated object detector (FOD). Yellow dots show fixation points, numbers in yellow fonts indicate the sequence of fixations and the bounding box is the final detection. Note that FOD does not have to fixate on the target object in order to localize it (example on the right).

model as the foveated visual field of our object detector (Figure 1). We call our detector as "the foveated object detector (FOD)" due to its foveated visual field.

The sizes of pooling regions in the visual field increase as a function of eccentricity from the fixation location. As the pooling regions get larger towards the periphery, more information is lost at these locations, which might seem to be a disadvantage, however, the exploration of the scene with the high resolution fovea through a guided search algorithm might mitigate the apparent loss of peripheral information. On the other hand, fewer computational resources are allocated to process these low resolution areas which, in turn, lower the computational cost. In this paper, we investigate the impact of using a foveated visual field on the detection performance and its computational cost savings.

1.1 Overview of our approach

The foveated object detector (FOD) mimics the process by which humans search for objects in scenes utilizing eye movements to point the high resolution fovea to points of interest (Figure 2). The FOD gets assigned an initial fixation point on the input image and collects information by extracting image features through its foveated visual field. The features extracted around the fixation point are at fine spatial scale while features extracted away from the fixation location at coarser scale. This fine-to-coarse transition is dictated by the pooling region sizes of the visual field. Then, based on the information collected, the FOD chooses the next fixation point and makes a saccade to that point. Finally, the FOD integrates information collected through multiple saccades and outputs object detection predictions.

Training such an object detector entails learning templates at all locations in the visual field. Because the visual field has varying resolution, the appearance of a target object varies depending on where it is located within the visual field. We use the HOG [6] as image features and a simplified version of the V1 model [15] to compute pooled features within the visual field. A mixture of linear templates is trained at selected locations in the visual field using a latent-SVM-like [14], [18] framework.

1.2 Contribution

We present an object detector that has a foveated visual field based on physiological measurements in primate visual cortex [15] and that models the appearance of target objects not only in the high resolution fovea but also in the periphery. . Importantly, the model is developed in the context of a modern object detection algorithm and a standard data-set (PASCAL VOC) allowing for the first time direct evaluation of the impact of a foveated visual system on an object detector.

We believe that object detection using a foveated visual field offers a novel and promising direction of research in the quest for an efficient alternative to the sliding window method, and also a possible explanation for why foveated visual systems might have evolved in organisms. We show that our method achieves greater computational savings than a state-of-the-art cascaded detection method. Another contribution of our work is the latent-LDA formulation (Section 2.4.2) where linear discriminant analysis is used within a latent-variable learning framework.

In the next section, we describe the FOD in detail and report experimental results in Section 3 which is followed by the related work section, conclusions and discussion.

2 The Foveated Object Detector (FOD)

2.1 Foveated visual field

The Freeman-Simoncelli (FS) model [15] is neuronal population model of V1 and V2 layers of the visual cortex.

The model specifies how responses are pooled (averaged together) hierarchically beginning from the lateral geniculate nucleus to V1 and then the V2 layer. V1 cells encode information about local orientation and spatial frequency whereas the cells in V2 pools V1 responses non-linearly to achieve selectivity for compound features such as corners and junctions. The model is based on findings and physiological measurements of the primate visual cortex and specifies the shapes and sizes of the receptive fields of the cells in V1 and V2. According to the model, the sizes of receptive fields increase linearly as a function of the distance from the fovea and this rate of increase in V2 is larger than that of V1, which means V2 pools larger areas of the visual field in the periphery. The reader is referred to [15] for further details.

We simplify the FS model in two ways. First, the model uses a Gabor filter bank to compute image features and we replace these with the HOG features [6], [14]. Second, we only use the V1 layer and leave the non-linear pooling at V2 as future work. We use this simplified FS model as the foveated visual field of our object detector which is shown in Figure 1. The fovea subtends a radius of 2 degrees. We also only simulate a visual field with a radius of 10 degrees which is sufficient to cover the test images presented at a typical viewing distance of 40 cm. The square boxes with white borders (Figure 1 represent the pooling regions within the fovea. The surrounding colored regions are the peripheral pooling regions. While the foveal regions have equal sizes, the peripheral regions grow in size as a function – which is specified by the FS model – of their distance to the center of the fovea. The color represents the weights that are used in pooling, i.e. weighted summation of, the underlying responses. A pooling region partly overlaps with its neighboring pooling regions (see the supplementary material of [15] for details). Assuming a viewing distance of 40cm, the whole visual field covers about a 500x500 pixel area (a pixel subtends $0.08°$). The foveal radius is 52 pixels subtending a visual angle of 4 degrees.

Given an image and a fixation point, we first compute the gradient at each pixel and then for each pooling region, the gradient magnitudes are pooled per orientation for the pixels that fall under the region. At the fovea, where the pooling regions are 8x8 pixels, we use the HOG features at the same spatial scale of the original DPM model [14], and in the periphery, each pooling region takes a weighted sum of HOG features of the 8x8 regions that are covered by that pooling region.

2.2 The model

The model M consists of a mixture of n components

$$M = \{(\mathbf{w}_i, \boldsymbol{\ell}_i) : i = 1, 2, \ldots, n\} \tag{1}$$

where $\mathbf{w}_i$ is a linear template and $\boldsymbol{\ell}_i$ is the location of the template with respect to the center of the visual field. The location variable $\boldsymbol{\ell}_i$ defines a unique bounding box within the visual field for the i^{th} template. Specifically, $\boldsymbol{\ell}_i = (\omega_i, h_i, x_i, y_i)$ is a 4-tuple whose variables respectively denote width, height and x,y coordinates of the i^{th} template within the visual field. The template, $\mathbf{w}_i$, is a matrix of weights on the features extracted from the pooling regions underlying the bounding box $\boldsymbol{\ell}_i$. The dimensionality of $\mathbf{w}_i$, i.e. the total number of weights, depends both on the width and height of its bounding box and its location in the visual field. A component within the fovea covers a larger number of pooling regions compared to a peripheral component with the same width and height, hence the dimensionality of a foveal template is larger. Three example components are illustrated in Figure 3 where the foveal component (red) covers 7x5 $= 35$ pooling regions while the (blue and green) peripheral components cover 15 and 2 regions, respectively. Since a fixed number of features[1] is extracted from each pooling region (regardless of its size), foveal components have higher-resolution templates associated with them.

2.2.1 Detection model

Suppose that we are given a model M that is already trained for a certain object class. The model is presented with an image I and assigned an initial fixation location $\boldsymbol{f}$. We are interested in searching for an object instance in I. Because the size of a searched object is not known apriori, the model has to analyze the input image at various scales. We use the same set of image scales given in [14] and use σ to denote a scale from that set. When used as a subscript to an image, e.g. I_σ, it denotes the scaled version of that image, i.e. width (and height) of I_σ is σ times the width (and height) of I. σ also applies to fixation locations and bounding boxes: if $\boldsymbol{f}$ denotes a fixation location $(\boldsymbol{f}_x, \boldsymbol{f}_y)$, then $\boldsymbol{f}_\sigma = (\sigma \boldsymbol{f}_x, \sigma \boldsymbol{f}_y)$; for a bounding box $\boldsymbol{b} = (w, h, x, y)$, $\boldsymbol{b}_\sigma = (\sigma w, \sigma h, \sigma x, \sigma y)$.

To check whether an arbitrary bounding box $\boldsymbol{b}$ within I contains an object instance, while the model is fixating at location f, we compute a detection score as

$$\mathrm{s}(I, \boldsymbol{b}, \boldsymbol{f}) = \max_{\sigma} \max_{c \in G(\boldsymbol{b}_\sigma, \boldsymbol{f}_\sigma)} \mathbf{w}^T \Psi(I_\sigma, \boldsymbol{f}_\sigma, c) \tag{2}$$

where $\Psi(I_\sigma, \boldsymbol{f}_\sigma, c)$ is a feature extraction function which returns the features of I_σ for component c (see Equation (1)) when the model is fixating at $\boldsymbol{f}_\sigma$. The vector $\mathbf{w}$ is the blockwise concatenation of the templates of all components. $\Psi(\cdot)$ effectively chooses which component to use, that is $\mathbf{w}^T \Psi(I_\sigma, \boldsymbol{f}_\sigma, c) = \mathbf{w}_c^T \Psi(I_\sigma, \boldsymbol{f}_\sigma, c)$. The fixation location ,$\boldsymbol{f}_\sigma$, together with the component c define a unique location, i.e. a bounding box, on I_σ. $G(\boldsymbol{b}_\sigma, \boldsymbol{f}_\sigma)$ returns the set of all components whose templates have a predetermined overlap (intersection over union should be at least 0.7 as in [14]) with $\boldsymbol{b}_\sigma$ when the model is fixating at $\boldsymbol{f}_\sigma$. During both training and testing, σ and c are latent variables for example $(I, \boldsymbol{b})$.

1. We use the feature extraction implementation of DPM (rel5) [14], [17], which extracts a 31-dimensional feature vector.

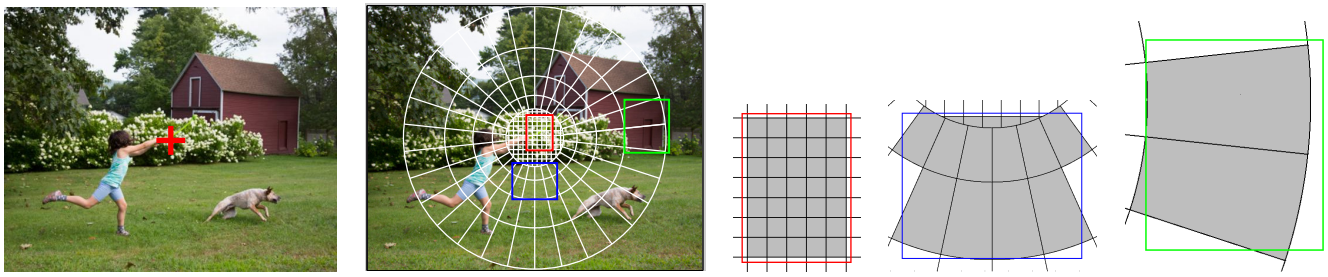

(a) (b) (c) (d) (e)

Fig. 3. Illustration of the visual field of the model. (a) The model is fixating at the red cross mark on the image. (b) Visual field (Figure 1) overlaid on the image, centered at the fixation location. White line delineate the borders of pooling regions. Nearby pooling regions do overlap. The weights (Figure 1) of a pooling region sharply decrease outside of its shown borders. White borders are actually iso-weight contours for neighboring regions. Colored bounding boxes show the templates of three components on the visual field: red, a template within the fovea; blue and green, two peripheral templates at 2.8 and 7 degree periphery, respectively. (c,d,e) Zoomed in versions of the red (foveal), blue (peripheral) and green (peripheral) templates. The weights of a template, $\mathbf{w}_i$, are defined on the gray shaded pooling regions.

Ideally, $s(I, \boldsymbol{b}, \boldsymbol{f}) > 0$ should hold for an appropriate $\boldsymbol{f}$ when I contains an object instance within $\boldsymbol{b}$. For an image that does not contain an object instance, $s(I, \boldsymbol{b} = \emptyset, \boldsymbol{f}) < 0$ should hold for any $\boldsymbol{f}$. For this to work, a subtlety in $G(\cdot)$'s definition is needed: $G(\emptyset, \boldsymbol{f})$ returns all components of the model (Equation (1)). During training (Section 2.4), this will enforce the responses of all components for a negative image to be suppressed down.

2.2.2 Integrating observations across multiple fixations

So far, we have looked at the situation where the model has made only one fixation. We describe in Section 2.3 how the model chooses the next fixation location. For now, suppose that the model has made m fixations, $\boldsymbol{f}_1, \boldsymbol{f}_2, \ldots, \boldsymbol{f}_m$, and we want to find out whether an arbitrary bounding box $\boldsymbol{b}$ contains an object instance. This computation involves integrating observations across multiple fixations, which is a considerably more complicated problem than the single fixation case. The Bayesian decision on whether $\boldsymbol{b}$ contains an object instance is based on the comparison of posterior probabilities:

$$\frac{P(y_{\boldsymbol{b}} = 1 | \boldsymbol{f}_1, \boldsymbol{f}_2, \ldots, \boldsymbol{f}_m, I)}{P(y_{\boldsymbol{b}} = 0 | \boldsymbol{f}_1, \boldsymbol{f}_2, \ldots, \boldsymbol{f}_m, I)} \gtrless 1 \tag{3}$$

where $y_{\boldsymbol{b}} = 1$ denotes the event that there is an object instance at location $\boldsymbol{b}$. We use the posteriors' ratio as a detection score, the higher it is the more likely $\boldsymbol{b}$ contains an instance. Computing the probabilities in (3) requires training a classifier per combination of fixation locations for each different value of m, which is intractable. We approximate it using a conditional independence assumption (derivation given in Appendix A):

$$\frac{P(y_{\boldsymbol{b}} = 1 | \boldsymbol{f}_1, \boldsymbol{f}_2, \ldots, \boldsymbol{f}_m, I)}{P(y_{\boldsymbol{b}} = 0 | \boldsymbol{f}_1, \boldsymbol{f}_2, \ldots, \boldsymbol{f}_m, I)} \approx \prod_{i=1}^{m} \frac{P(y_{\boldsymbol{b}} = 1 | \boldsymbol{f}_i, I)}{P(y_{\boldsymbol{b}} = 0 | \boldsymbol{f}_i, I)}. \tag{4}$$

We model the probability $P(y_{\boldsymbol{b}} = 1 | \boldsymbol{f}, I)$ using a classifier and use the sigmoid transfer function to convert raw classification scores to probabilities:

$$P(y_{\boldsymbol{b}} = 1 | \boldsymbol{f}, I) = \frac{1}{1 + e^{(-s(I, \boldsymbol{b}, \boldsymbol{f}))}}. \tag{5}$$

We simplify the computation in (4) by taking the log (derivation given in Appendix B):

$$\log \left(\prod_{i=1}^{m} \frac{P(y_{\boldsymbol{b}} = 1 | f_i, I)}{P(y_{\boldsymbol{b}} = 0 | f_i, I)} \right) = \sum_{i=1}^{m} s(I, \boldsymbol{b}, \boldsymbol{f}_i). \tag{6}$$

Taking the logarithm of posterior ratios does not alter the ranking of detection scores for different locations, i.e. $\boldsymbol{b}$'s, because logarithm is a monotonic function. In short, the detection score computed by the FOD for a certain location $\boldsymbol{b}$, is the sum of the individual scores for $\boldsymbol{b}$ computed at each fixation.

After evaluating (6) for a set of candidate locations, final bounding box predictions are obtained by non-maxima suppression [14], i.e. given multiple predictions for a certain location, all predictions except the one with the maximal score are discarded.

2.3 Eye movement strategy

We use the maximum-a-posteriori (MAP) model [4] as the basic eye movement strategy of the FOD. The MAP model is shown to be consistent with human eye movements in a variety of visual search tasks [4], [47]. Studies have demonstrated that in some circumstances human saccade statistics better match an ideal searcher [31] that makes eye movements to locations that maximize the accuracy of localizing targets, yet in many circumstances the MAP model approximates the ideal searcher [32], [51] but is computationally more tractable for objects in real scenes. The MAP model select the location with the highest posterior probability of containing the target

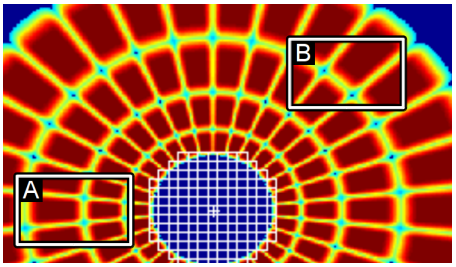

Fig. 4. Two bounding boxes (A,B) are shown on the visual field. While box A covers a large portion of the pooling regions that it intersects with, box B's coverage is not as good. Box B is discarded as it does not meet the overlap criteria (see text), therefore a component for B in the model is not created.

object as the next fixation location, that is $\boldsymbol{f}_{i+1}$ = center of $\boldsymbol{\ell}^*$ where

$$\boldsymbol{\ell}^* = \arg\max_{\boldsymbol{\ell}} P(y_{\boldsymbol{\ell}} = 1 | \boldsymbol{f}_1, \boldsymbol{f}_2, \ldots, \boldsymbol{f}_i, I). \tag{7}$$

Finding the maximum of the posterior above is equivalent to finding the maximum of the posterior ratios,

$$\arg\max_{\boldsymbol{\ell}} P(y_{\boldsymbol{\ell}} = 1 | \boldsymbol{f}_1, \ldots, \boldsymbol{f}_i, I) = \arg\max_{\boldsymbol{\ell}} \frac{P(y_{\boldsymbol{\ell}}=1|\boldsymbol{f}_1,\ldots,\boldsymbol{f}_i,I)}{P(y_{\boldsymbol{\ell}}=0|\boldsymbol{f}_1,\ldots,\boldsymbol{f}_i,I)} \tag{8}$$

since for two arbitrary locations $\boldsymbol{\ell}_1, \boldsymbol{\ell}_2$; let $p_1 = P(y_{\boldsymbol{\ell}_1} = 1|\cdot)$ and $p_2 = P(y_{\boldsymbol{\ell}_2} = 1|\cdot)$, then we have

$$\frac{p_1}{1-p_1} > \frac{p_2}{1-p_2} \implies p_1 > p_2. \tag{9}$$

2.4 Training the model

2.4.1 Initialization

A set of dimensions (width and height) is determined from the bounding box statistics of the examples in the training set as done in the initialization of the DPM model [14]. Then, for each width and height, new components with these dimensions are created to tile the entire visual field. However, the density of components in the visual field is not uniform. Locations, i.e. bounding boxes, that do not overlap well with the underlying pooling regions are discarded. To define goodness of overlap, a bounding box is said to intersect with an underlying pooling region if more than one fifth of that region is covered by the bounding box. Overlap is the average coverage across the intersected regions. If the overlap is more than 75%, then a component for that location is created, otherwise the location is discarded (see Figure 4 for an example). In addition, no components are created for locations that are outside of the visual field. Weights of the component templates ($\mathbf{w}_i$) are initialized to arbitrary values. Training the model is essentially optimizing these weights on a given dataset.

2.4.2 Training

Consider a training set $\mathcal{D} = \{(I_i, \boldsymbol{b}_i)\}_{i=1}^K$ where I_i is an image and $\boldsymbol{b}_i$ a bounding box and K is the total number of examples. If I_i does not contain any positive examples, i.e. object instances, then $\boldsymbol{b}_i = \emptyset$. Following the DPM model [14], we train model templates using a latent-SVM formulation:

$$\arg\min_{\mathbf{w}} \frac{1}{2}||\mathbf{w}||_2^2 + C\sum_{i=1}^{K} \sum_{\boldsymbol{f} \in F(I_i, \boldsymbol{b}_i)} \max(0, 1 - y_i s(I_i, \boldsymbol{b}_i, \boldsymbol{f})). \tag{10}$$

where $y_i = 1$ if $\boldsymbol{b}_i \neq \emptyset$ and $y_i = -1$, otherwise. The set $F(I_i, \boldsymbol{b}_i)$ denotes the set of all *feasible* fixation locations for example $(I_i, \boldsymbol{b}_i)$. For $\boldsymbol{b}_i \neq \emptyset$, a fixation location is considered feasible if there exists a model component whose bounding box overlaps with $\boldsymbol{b}_i$. For $\boldsymbol{b}_i = \emptyset$, all possible fixation locations on I_i are considered feasible.

Optimizing the cost function in (10) is manageable for mixtures with few components, however, the FOD has a large number of components in its visual field (typically, for an object class in the PASCAL VOC 2007 dataset [12], there are around 500-700) and optimizing this cost function becomes prohibitive in terms of computational cost. As an alternative, cheaper linear classifiers can be used. Recently, linear discriminant analysis (LDA) has been used in object detection ([18]) producing surprisingly good results with much faster training time. Training a LDA classifier amounts to computing $\Sigma^{-1}(\mu_1 - \mu_0)$ where μ_1 is the mean of the feature vectors of the positive examples, μ_0 is the same for the negative examples and Σ is the covariance matrix of these features. Here, the most expensive computation is the estimation of Σ, which is required for each template with different dimensions. However, it is possible to estimate a global Σ from which covariance matrices for templates of different dimensions can be obtained [18]. For the FOD, we estimate the covariance matrices for the foveal templates and estimate the covariance matrices for peripheral templates by applying the feature pooling transformations to the foveal covariance matrices.

We propose to use LDA in a latent-SVM-like framework as an alternative to the method in [18] where positive examples are clustered first and then a LDA classifier is trained per cluster. Consider the t^{th} template, $\mathbf{w}_t$. LDA gives us that LDA gives us that $\mathbf{w}_{t,\text{LDA}} = \Sigma_t^{-1}(\mu_t^{\text{pos}} - \mu_t^{\text{neg}})$ where Σ_t is the covariance matrix for template t, μ_t^{pos} and μ_t^{neg} are the mean of positive and negative feature vectors, respectively, assigned to template t. We propose to apply an affine transformation to the LDA classifier:

$$\mathbf{w}_t = \begin{bmatrix} \alpha_t & & 0 \\ & \alpha_t & \\ & \ddots & \\ 0 & & \alpha_t \\ & & \beta_t \end{bmatrix} \begin{bmatrix} \mathbf{w}_{t,\text{LDA}} \\ 1 \end{bmatrix} = \begin{bmatrix} \alpha_t \mathbf{w}_{t,\text{LDA}} \\ \beta_t \end{bmatrix} \tag{11}$$

and modify the cost function as

$$\arg\min_{\boldsymbol{\alpha},\boldsymbol{\beta}} \Bigg(\tfrac{1}{2}||\mathbf{w}||_2^2 + C\sum_{t=1}^{N} \max(0, 1 + {\mathbf{w_t}}^T \mu_t^{neg}) + C\sum_{i\in\{i|\boldsymbol{b}_i\neq\emptyset\}} \sum_{\boldsymbol{f}\in F(I_i,\boldsymbol{b}_i)} \max(0, 1 - y_i s(I_i, \boldsymbol{b}_i, \boldsymbol{f})) \Bigg) \quad (12)$$

where the first summation pushes the score of the mean of the negative examples to under zero and the second summation, taken over positive examples only, pushes the scores to above 0. $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ are appropriate blockwise concatenation of α_t and β_ts. C is the regularization constant. Overall, this optimization effectively calibrates the dynamic ranges of different templates' responses in the model so that the scores of positive examples and negative means are pushed away from each other while the norm of $\mathbf{w}$ is constraint to prevent overfitting. This formulation does not require the costly mining of hard-negative examples of latent-SVM. We call this formulation (Equation (12)) as latent-LDA.

To optimize (12), we use the classical coordinate-descent procedure. We start by initializing $\mathbf{w}$ by training on warped-positive examples as in [14]. Then, we alternate between choosing the best values for the latent variables while keeping $\mathbf{w}$ fixed, and optimizing for $\mathbf{w}$ while keeping the latent variables of positive examples fixed.

3 Experiments

We evaluated our method on the PASCAL VOC 2007 detection (`comp3`) challenge dataset and protocol (see [12] for details). All results are obtained by training on the `train+val` split and testing on the `test` split.

3.1 Comparison of SW based methods

We first compared our SW implementation, which corresponds to using foveal templates only, to three state-of-the-art methods that are also SW based [14], [18], [29]. Table 1 gives the AP (average precision) results, i.e. area under the precision-recall curve per class, and mean AP (mAP) over all classes. Originally, the deformable parts model (DPM) uses object parts, however, in order to make a fair comparison with our model, we disabled its parts. The first row of Table 1 shows the latest version of the DPM system [17] with the parts-learning code disabled. The second row shows results for another popular SVM-based system, known as the exemplar-SVM (E-SVM), which also only models whole objects, not its parts. Finally, the third row shows results from a LDA-based system, "discriminative decorrelation for classification" (DCC) [18]. All three systems are based on HOG features and mixture of linear templates. The results show that SVM based systems perform better than the LDA based systems, which is not a surprising finding since it is well known that discriminative models outperform generative models in classification tasks. However, LDA's advantage against this performance loss is that it is ultra fast to train, which is exactly the reason we chose to use LDA instead of SVM. Once the background covariance matrices are estimated (which can be done once and for all [18]), training is as easy as taking the average of the feature vectors of positive examples and doing a matrix multiplication. We estimated the time that training a SVM based system for our FOD to be about 300 hours (approximately 2 weeks) for a single object class, whereas the LDA based system can be trained under an hour on the same machine which has an Intel i7 processor.

Although our SW method achieves the same mean AP (mAP) score as the DCC method [18], the latter has a detection model with higher computational cost. We use 2 templates per class while DCC trains more than 15 templates per class within an exemplar-SVM [29]-like framework. DCC considers the dot product of the feature vector of the detection window with every exemplar within a cluster, which basically means that a detection window is compared to all positive examples in the training set. In our case, the number of dot products considered per detection window is equal to the number of templates, which is 2 in this paper, which clearly demonstrates the advantage of our latent-LDA approach over DCC [18].

3.2 Comparison of FOD with SW

Next, we compared the performance of FOD with our SW method. We experimented with two eye movement strategies, MAP (Section 2.3) and random strategy to demonstrate the importance of guidance of eye movements.

Table 2 shows the AP scores for FOD with different eye movement strategies and different number of fixations. We also include in this table the "Our SW" result from Table 1 for ease of reference. The MAP and random strategies are denoted with MAP and RAND, respectively. Because the model accuracy results will depend on initial point of fixation, we ran the models with different initial points of fixation. The presence of a suffix on a model refers to the location of the initial fixation: "-C" stands for the center of the input image, i.e. $(0.5, 0.5)$ in normalized image coordinates where the top-left corner is taken as $(0, 0)$ and the bottom-right corner is $(1, 1)$; and "-E" for the two locations at the left and right edges of the image, 10% of the image width away from the image border, that is $(0.1, 0.5)$ and $(0.9, 0.5)$. MAP-E and RAND-E results are the performance average of two different runs, one with initial fixation close to the left edge of the image, the other run close to the right edge of the image. For the random eye movement, we report the 95% confidence interval for AP over 10 different runs. We ran all systems for a total of 5 fixations. Table 2 shows results for after 1,3 and 5 fixations. A condition with one fixation is a model that makes decisions based only on the initial fixation.

TABLE 1
Average precision (AP) scores of SW based methods on the PASCAL VOC 2007 dataset.

	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	person	plant	sheep	sofa	train	tv	mAP
DPM [14]	23.6	48.6	9.7	11.0	19.3	40.4	45.2	12.4	15.4	19.4	17.4	4.0	44.7	36.4	31.2	10.9	14.1	19.5	32.2	37.0	24.6
E-SVM [29]	20.4	40.7	9.3	10.0	10.3	31.0	40.1	9.6	10.4	14.7	2.3	9.7	38.4	32.0	19.2	9.6	16.7	11.0	29.1	31.5	19.8
DCC [18]	17.4	35.5	9.7	10.9	15.4	17.2	40.3	10.6	10.3	14.3	4.1	1.8	39.7	26.0	23.1	4.9	14.1	8.7	22.1	15.2	17.1
Our SW	17.5	28.6	9.7	10.4	17.3	29.8	36.7	7.9	11.2	21.0	2.3	2.7	30.9	21.1	19.7	3.0	9.2	13.7	23.5	25.2	17.1

TABLE 2
AP scores and relative computational costs of SW and FOD on the PASCAL VOC 2007 dataset.

		aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	person	plant	sheep	sofa	train	tv	mAP	Comp. Cost
Our SW		17.5	28.6	9.7	10.4	17.3	29.8	36.7	7.9	11.2	21.0	2.3	2.7	30.9	21.1	19.7	3.0	9.2	13.7	23.5	25.2	17.1	100
MAP-C	**1**	17.0	21.1	4.9	9.8	9.3	27.4	27.9	8.5	3.7	12.8	2.0	4.3	29.7	19.7	18.2	1.2	10.7	14.0	26.2	21.8	14.5	11.5
	3	17.4	27.7	10.1	10.6	10.4	30.8	31.6	8.4	10.4	17.2	2.1	3.4	33.3	21.1	18.7	3.4	7.6	15.4	26.4	23.5	16.5	31.2
	5	17.0	28.6	10.0	10.7	11.2	31.0	34.0	8.3	10.6	18.2	2.1	3.4	34.2	21.8	19.7	2.8	8.1	15.1	27.8	24.0	16.9	49.6
MAP-E	**1**	1.6	7.1	4.1	5.6	9.1	8.7	11.7	6.0	3.6	10.2	2.0	2.2	8.5	10.2	13.5	1.3	6.8	8.0	10.6	10.3	7.1	8.7
	3	13.0	24.6	9.9	9.8	10.7	27.2	29.3	7.4	10.4	16.4	3.7	2.2	30.6	20.8	16.9	3.3	11.2	13.8	23.0	24.1	15.4	28.1
	5	15.1	28.0	9.9	10.4	11.6	29.9	33.0	8.3	10.6	18.7	2.7	4.1	33.7	22.6	18.9	3.1	7.1	14.7	25.5	25.2	16.7	46.9
RAND	**1**	8.2	9.3	5.5	9.3	7.8	12.2	16.2	6.1	6.8	7.5	1.6	2.5	10.6	9.1	9.9	1.9	5.0	6.7	11.2	10.0	7.9±1.4	similar to above
	3	9.6	13.0	3.2	9.6	9.3	16.9	23.5	8.8	9.4	9.9	1.8	3.2	16.5	12.3	12.2	2.7	3.9	9.3	16.9	11.7	10.2±0.9	
	5	10.9	15.3	3.8	9.7	9.6	20.5	26.3	9.3	9.5	10.6	1.5	3.1	20.9	13.7	13.5	2.7	3.9	12.0	18.9	12.4	11.4±1.0	
RAND-C	**1**	This row is the same with the "MAP-C, 1" above.																					
	3	17.5	20.4	3.7	10.0	9.3	28.6	27.4	11.5	6.7	11.8	1.7	3.5	31.7	18.0	15.4	2.7	5.4	15.2	26.1	15.8	14.1±0.5	"
	5	17.6	21.4	5.2	9.9	9.7	28.1	28.6	11.4	9.6	12.1	1.6	3.5	30.0	17.9	15.3	3.7	6.7	14.4	25.4	15.9	14.4±0.7	
RAND-E	**1**	This row is the same with the "MAP-E, 1" above.																					
	3	9.1	13.1	2.8	9.7	9.4	17.8	22.5	9.0	6.6	10.7	2.3	3.7	14.9	12.0	14.9	1.3	3.9	2.4	13.6	14.1	9.7±0.7	"
	5	10.7	15.9	4.1	8.7	9.5	21.9	26.0	8.2	9.7	11.6	1.7	4.3	17.6	13.7	14.1	1.9	5.7	4.8	15.7	15.8	11.1±1.1	

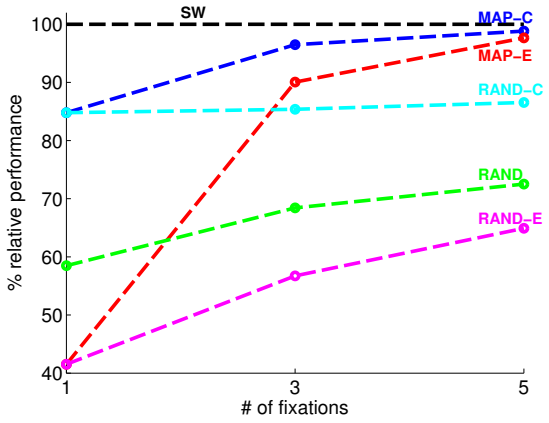

Fig. 5. Ratio of mean AP scores of FOD systems relative to that of the SW system. Graph shows two eye movement algorithms: maximum aposteriori probability (MAP) and random (RAND) and two starting points (C: center; E: edge).

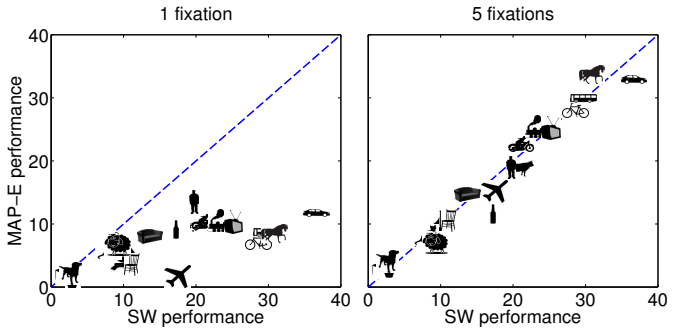

Fig. 6. AP scores achieved by SW and MAP-E per class.

The results show that the FOD using the MAP rule with 5 fixations (MAP-C,5 for short) performs nearly as good as the SW (a difference of 0.2 in mean AP).

Figure 5 shows the ratio of mean AP for the FOD with the various eye movement strategies to that of the SW system (relative performance) as a function of fixation. The relative performance of the MAP-C to SW (AP of MAP-C divided by AP of SW) is 98.8% for 5 fixations, 96.5% for 3 fixations and 84.8% for 1 fixation. The FOD with eye movement guidance towards the target (MAP-C,5) achieves or exceeds SW's performance with only 1 fixation in 4 classes, with 3 fixations in 7 classes, with 5 fixations in 2 classes. For the remaining of 7 classes, FOD needs more than 5 fixations to achieve SW's performance.

MAP-C performs quite well (84.8% relative performance) even with 1 fixation. The reason behind this result is the fact that, on average, bounding boxes in the PASCAL dataset cover a large portion of the images (average bounding box area normalized by image area is 0.2) and are located at and around the center [44]. To reduce the effects of these biases about the location of object placement on the results, we assessed the models with an initial fixation close to the edge of the image (MAP-E). When the initial fixation is closer to the edge of the image, performance is initially worse than when the initial fixation is at the center of the image, The difference in performance diminishes achieving similar performance with five fixations (0.2 difference in mean

AP). Figure 6 shows how the distribution of AP scores for different object classes for MAP-E improves from 1 fixation to 5 fixations

3.2.1 Importance of the guidance algorithm

To assess the importance of guided saccades towards the target we compared performance of the MAP model against FOD that guides eye movements based on a random eye movement generator.

Figure 5 allows comparisons of the relative performance of the MAP FOD and those with a random eye movement strategy. The performance gap between MAP-C, RAND-C pair and MAP-E,RAND-E pair shows that MAP eye movement strategy is effective in improving the performance of the system.

3.3 Computational cost

The computational complexity of the SW method is easily expressed in terms of image size. However, this is not the case for our model. The computational complexity of FOD is $O(mn)$ where m is the number of fixations and n is the total number of components, hence templates, on the visual field. These numbers do not explicitly depend on the image size; so in this sense, the complexity of FOD is $O(1)$ in terms of image size. Currently, m is given as an input parameter but if it were to be automated, e.g. to achieve a certain detection accuracy, m would implicitly depend on several factors such as the difficulty of the object class, the location and size distribution of positive examples. Targets that are small (relative to the image size) and that are located far away from the initial fixation location would require more fixations to get a certain detection accuracy. The number of components, n, depends on both the visual field parameters (number of angle and eccentricity bins which, in our case, are fixed based on the Freeman-Simoncelli model [15]) and the bounding box statistics of the target object. These dependencies make it difficult to express the theoretical complexity in terms of input image size. For this reason, we compare the computational costs of FOD and SW in a practical framework, expressed in terms of the total number of operations performed in template evaluations.

In both SW based methods and the FOD, linear template evaluations, i.e. taking dot-products, is the main time consuming operation. We define the computational cost of a method based on the total number of template evaluations it executes (as also done in [46]). A model may have several templates with different sizes, so instead of counting each template evaluation as 1 operation, we take into account the dimensionalities of the templates. For example, the cost of evaluating a (6-region)x(8-region) HOG template is counted as 48 operations.

It is straightforward to compute the computational cost (as defined above) of the SW method. For the FOD, we run the model on a subset of the testing set and count

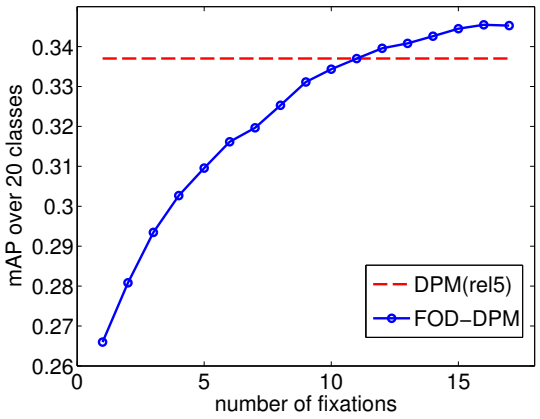

Fig. 7. FOD-DPM's performance (mean AP over 20 classes) as a function of number of fixations. FOD-DPM achieves DPM's performance at 11 fixations and exceeds it with more fixations.

the number of operations actually performed. Note that, in order to compute a detection score, the FOD first performs a feature pooling (based on the location of the component in the visual field) and then a linear template evaluation. Since these are both linear operations, we combine them into a single linear template. The last column of Table 2 gives the computational costs of the SW method and the FOD. For the FOD the computational cost is reported as a function of different number of fixations. For ease of comparison, we normalized the costs so that the SW method performs 100 operations in total. The results show that FOD is computationally more efficient than SW. FOD achieves 98.8% of SW's performance at 49.6% of the computational cost of SW. Note that this saving is not directly comparable to that of the cascaded detection method reported in [13] because FOD's computational savings comes about from fewer root filter evaluations, whereas in [13] a richer model (DPM, root filters and part filters) is used and the savings are associated to fewer evaluations in the part filters (i.e., the model applies the root filters at all locations first and sequentially running other filters on the non-rejected locations).

3.4 Using richer models to increase performance

To directly compare the computational savings of the FOD model to a cascade-type object detector, we used a richer and more expensive detection model at the fovea. This is analogous to the cascaded detection idea where cheaper detectors are applied first and more expensive detectors are applied later on the locations not rejected by the cheaper detectors. To this end, we run our FOD and after each fixation we evaluate the full DPM detector (root and part filters together) [17] only at foveal locations that score above a threshold which is determined on the training set to achieve high recall rate (95%). We call this approach "FOD-DPM cascade" or FOD-DPM for short. Table 3 and Figure 7 give the

TABLE 3
AP scores and relative computational costs of FOD-DPM and DPM on the PASCAL VOC 2007 dataset.

		aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	person	plant	sheep	sofa	train	tv	mAP	Comp. Cost
DPM(rel5)		33.2	60.3	10.2	16.1	27.3	54.3	58.2	23.0	20.0	24.1	26.7	12.7	58.1	48.2	43.2	12.0	21.1	36.1	46.0	43.5	33.7	100
FOD-DPM	**1**	31.0	37.1	10.0	14.3	12.9	47.1	46.7	28.0	9.3	15.5	26.2	10.7	56.0	39.7	29.4	9.8	15.5	27.6	43.4	21.5	26.6	0.46
	5	32.3	50.0	9.8	15.2	21.8	50.0	63.0	25.9	17.1	20.5	25.4	9.7	61.4	44.6	38.0	9.2	19.7	30.1	43.1	32.1	31.0	1.84
	9	33.2	56.6	9.9	15.6	25.3	54.6	65.3	25.3	19.8	22.0	24.9	9.4	60.9	50.8	41.7	10.0	20.4	34.9	44.3	37.3	33.1	3.09
	13	**33.4**	59.9	10.0	15.7	27.2	**54.8**	**65.7**	**25.0**	**20.5**	22.0	24.8	9.2	**62.0**	**51.9**	**44.5**	10.2	20.9	**36.8**	**46.2**	40.9	**34.1**	4.16

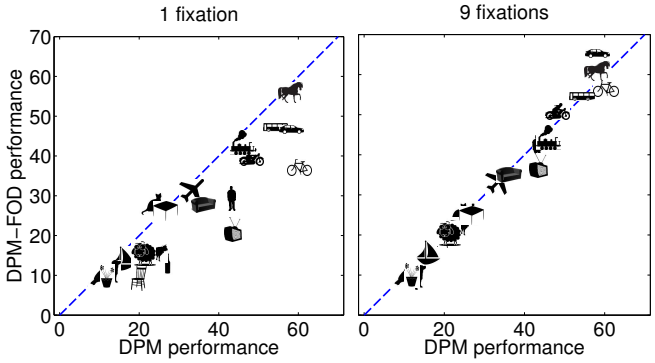

Fig. 8. AP scores achieved by FOD-DPM and DPM per class.

performance result of this approach. FOD-DPM achieves a similar average performance to that of DPM (98.2% relative performance, 0.6 AP gap) using 9 fixations and exceeds DPM's performance starting from 11 fixations. On some classes (e.g. bus, car, horse), FOD-DPM exceeds DPM's performance probably due to lesser number of evaluations and reduced false positives; on other cases (e.g. bike, dog, tv) FOD-DPM underperforms probably due to low recall rate of the FOD detector for these classes. Figure 8 gives per class AP scores of FOD-DPM and DPM to demonstrate the improvement from 1 to 9 fixations.

We compare the computational complexities of FOD-DPM and DPM by their total number of operations as defined above. For a given object class, DPM model has 3 root filters and 8 6x6 part filters. It is straightforward to calculate the number of operations performed by DPM as it uses the SW method. For FOD-DPM, the total number of operations is calculated by adding: 1) FOD's operations and 2) DPM's operations at each high-scoring foveal detection $\boldsymbol{b}$, one DPM root filter (with the most similar shape as $\boldsymbol{b}$) and 8 parts evaluated at all locations within the boundaries of this root filter. Note that we ignore the time for optimal placing of parts in both DPM and FOD-DPM. Cost of feature extraction is also not included as the two methods use the same feature extraction code. We report the computational costs of FOD-DPM and DPM in the last column of Table 3. The costs are normalized so that DPM's cost is 100 operations. Results show that FOD-DPM drastically reduces the cost from 100 to 3.09 for 9 fixations. Assuming both methods are implemented equally efficiently, this would translate to an approximately 32x speed-up which is better than the 20x speed-up reported for a cascaded evaluation of DPM [13]. These results demonstrate the effectiveness of our foveated object detector in guiding the visual search.

Finally, in Figure 9 we give sample detections by the FOD system. We ran the trained bicycle, person and car models on an image outside of the PASCAL datasaet. The models were assigned the same initial location and we ran them for 3 fixations. Results show that the each model fixates at different locations, and these locations are attracted towards instances of the target objects being searched.

4 Related Work

The sliding window (SW) method is the dominant model of search in object detection. The complexity of identifying object instances in a given image is $\mathcal{O}(mn)$ where m is the number of locations to be evaluated and n is the number of object classes to be searched for. Efficient alternatives to sliding windows can be categorized in two groups: (i) methods aimed at reducing m, (ii) methods aimed at reducing n. Since typically $m >> n$, the are a larger number efforts in trying to reduce m, however, reducing the contribution of the number of object classes has recently been receiving increasing interest as search for hundreds of thousands of object classes has started to be tackled [7]. According to this categorization, our proposed FOD method falls into the first group as it is designed to locate object instances by making a set of sequential fixations where in each fixation only a sparse set of locations are evaluated.

4.1 Reducing the number of evaluated locations (m)

In efforts to reduce the number of locations to be evaluated, one line of research is the branch-and-bound methods ([20], [24]) where an upper bound on the quality function of the detection model is used in a global branch and bound optimization scheme. Although the authors provide efficiently computable upper bounds for popular quality functions (e.g. linear template, bag-of-words, spatial pyramid), it might not be trivial to derive suitable upper bounds for a custom quality function. Our method, on the other hand, uses binary classification detection model and is agnostic to the quality function used.

Another line of research is the casdaded detection framework ([13], [23], [48]) where a series of cheap to expensive tests are done to locate the object. Cascaded detection is similar to our method in the sense that

Fig. 9. Fixation locations and bounding box predictions of FOD for different object classes (bicycle, person, and car from left to right) but for the same image and initial point of fixation.

simple, coarse and cheap evaluations are used together with complex, fine and expensive evaluations. However, we differ with it in that it is essentially a sliding window method with a coarse-to-fine heuristic used to reduce the number of total evaluations. Another coarse-to-fine search scheme is presented in [34] where a set of low to high resolution templates are used. The method starts by evaluating the lowest resolution template – which is essentially a sliding window operation – and selecting the high responding locations for further processing with higher resolution templates. Our method, too, uses a set of varying resolution templates; however, these templates are evaluated at every fixation instead of serializing their evaluations with respect to resolution.

In [46], a segmentation based method is proposed to yield a small set of locations that are likely to corresponds to objects, which are subsequently used to guide the search in a selective manner. The locations are identified in an object class-independent way using an unsupervised multiscale segmentation approach. Thus, the method evaluates the same set of locations regardless of which object class is being searched for. In contrast, in our method, selection of locations to be foveated is guided by learned object class templates.

The method in [1], similar to ours, works like a fixational system: at a given time step, the location to be evaluated next is decided based on previous observations. However, there are important differences. In [1], only a single location is evaluated at a time step whereas we evaluate all template locations within the visual field at each fixation. Their method returns only one box as the result whereas our method is able to output many predictions.

There are also vector quantization based methods [19], [21], [40] aiming to reduce the time required to compute linear template evaluations. These methods to reduce the contribution of m in $\mathcal{O}(mn)$ are orthogonal to our foveated approach. Thus, vector quantization approaches can be integrated with the proposed foveated object detection method.

4.2 Reducing the number of evaluations of object classes(n)

Works in this group aim to reduce the time complexity contributed by the number of object classes. The method proposed in [7] accelerates the search by replacing the costly linear convolution by a locality sensitive hashing scheme that works on non-linearly coded features. Although they evaluate all locations in a given image, their approach scale constantly over the number of classes, which enables them to evaluate thousands of object classes in a very short amount of time.

Another method [42] uses a sparse representation of object part templates, and then uses the basis of this representation to reconstruct template responses. When the number of object categories is large, sparse representation serves as a shared dictionary of parts and accelerates the search.

Another line of research (e.g. [3], [16], [36]) accelerate search by constructing classifier hierarchies. These generally work by pruning unlikely classes while descending the classifier hierarchy.

Importantly, the way the methods in this group accelerate search is orthogonal to the savings proposed by using a foveated visual field. Therefore, these methods are complementary and can be integrated with our method to further accelerate search.

In the context of the references listed in this and the previous sections, our method of search through fixations using a non-uniform foveated visual field is novel.

4.3 Biologically inspired methods

There have been previous efforts, (e.g. [41]), on biologically inspired object recognition. However, these models do not have a foveated visual field and thus do not execute eye movements. More recent work has implemented biologically inspired search methods. In [11], a fixed, pre-attentive, low-resolution wide-field camera is combined with a shiftable, attentive, high-resolution narrow-field camera, where the pre-attentive camera generates saccadic targets for the attentive, high-resolution camera. The fundamental difference between this and our method is that while their pre-attentive system has the same coarse resolution everywhere in

the visual field, our method, which is a model of the V1 layer of the visual cortex, has a varying resolution that depends on the radial distance to the center of the fovea. There have been previous efforts to create foveated search models with eye movements [30], [31], [38], [51]. Such models have been applied mostly to detect simple signals in computer generated noise [31], [51] and used as benchmarks to compare against human eye movements and performance.

Other biologically inspired methods include the target acquisition model (TAM) [50], [52], the Infomax model [5] and artificial neural network based models [2], [25]. TAM is a foveated model and it uses scale invariant feature transform (SIFT) features [27] for representation and utilizes a training set of images to learn the appearance of the target object. However, it does not include the variability in object appearance due to scale and viewpoint, and the evaluation is done by placing the objects on a uniform background. The Infomax, on the other hand, can use any previously trained object detector and works on natural images. They report successful results on a face detection task. Both TAM and Infomax uses the same template for all locations in the visual field while our method uses different templates for different locations. [25] was applied to image categorization and [2] to object tracking in videos. Critically, none of these models have been tested on standard object detection datasets nor they have been compared to a SW approach to evaluate the potential performance loss and computational savings of modeling a foveated visual field.

5 Conclusions and Discussion

We present an implementation of a foveated object detector with a recent neurobiologically plausible model of pooling in the visual periphery and report the first ever evaluation of a foveated object detection model on a standard data set in computer vision (PASCAL VOC 2007). Our results show that the foveated method achieves nearly the same performance as the sliding window method at 49.6% of sliding window's computational cost. Using a richer model (such as DPM [14]) to evaluate high-scoring locations, FOD is able to outperform the DPM with more computational savings than a state-of-the-art cascaded detection system [13]. These results suggest that using a foveated visual system offers a promising potential for the development of more efficient object detectors.

Appendix A
Approximation of the Bayesian decision

Derivation for Equation (4):

$$\frac{P(y_{\boldsymbol{b}}=1|\boldsymbol{f}_1,\ldots,\boldsymbol{f}_m,I)}{P(y_{\boldsymbol{b}}=0|\boldsymbol{f}_1,\ldots,\boldsymbol{f}_m,I)} = \frac{P(f_1,\ldots,\boldsymbol{f}_m|y_{\boldsymbol{b}}=1,I)P(y_{\boldsymbol{b}}=1|I)}{P(f_1,\ldots,\boldsymbol{f}_m|y_{\boldsymbol{b}}=0,I)P(y_{\boldsymbol{b}}=0|I)} \tag{13}$$

$$\approx \prod_{i=1}^{m} \frac{P(\boldsymbol{f}_i|y_{\boldsymbol{b}}=1,I)}{P(\boldsymbol{f}_i|y_{\boldsymbol{b}}=0,I)} \frac{P(y_{\boldsymbol{b}}=1|I)}{P(y_{\boldsymbol{b}}=0|I)} = \prod_{i=1}^{m} \frac{P(y_{\boldsymbol{b}}=1|\boldsymbol{f}_i,I)}{P(y_{\boldsymbol{b}}=0|\boldsymbol{f}_i,I)}. \tag{14}$$

Appendix B
Detection score after multiple fixations

Derivation for Equation (6):

$$\log\left(\prod_{i=1}^{m} \frac{P(y_{\boldsymbol{b}}=1|\boldsymbol{f}_i,I)}{P(y_{\boldsymbol{b}}=0|\boldsymbol{f}_i,I)}\right) = \sum_{i=1}^{m} \log\left(\frac{P(y_{\boldsymbol{b}}=1|\boldsymbol{f}_i,I)}{1-P(y_{\boldsymbol{b}}=1|\boldsymbol{f}_i,I)}\right) \tag{15}$$

using (5), we get

$$\sum_{i=1}^{m} \log\left(\frac{P(y_{\boldsymbol{b}}=1|\boldsymbol{f}_i,I)}{1-P(y_{\boldsymbol{b}}=1|\boldsymbol{f}_i,I)}\right) = \sum_{i=1}^{m} \log\left(\frac{\frac{1}{1+e^{-s(I,\boldsymbol{b},\boldsymbol{f}_i)}}}{1-\frac{1}{1+e^{-s(I,\boldsymbol{b},\boldsymbol{f}_i)}}}\right) \tag{16}$$

$$= \sum_{i=1}^{m} \log\left(\frac{1}{e^{-s(I,\boldsymbol{b},\boldsymbol{f}_i)}}\right) = \sum_{i=1}^{m} s(I,\boldsymbol{b},\boldsymbol{f}_i). \tag{17}$$

References

[1] B. Alexe, N. Heess, Y. W. Teh, and V. Ferrari. Searching for objects driven by context. In *Advances in Neural Information Processing*, pages 1–9, 2012.

[2] L. Bazzani, N. de Freitas, H. Larochelle, V. Murino, and J.-A. Ting. Learning attentional policies for tracking and recognition in video with deep networks. In *Int'l Conf. on Machine Learning*, 2011.

[3] S. Bengio, J. Weston, and D. Grangier. Label Embedding Trees for Large Multi-Class Tasks. In J. Lafferty, C. Williams, J. Shawe-Taylor, R. Zemel, and A. Culotta, editors, *Advances in Neural Information Processing*, page 163171. 2010.

[4] B. R. Beutter, M. P. Eckstein, and L. S. Stone. Saccadic and perceptual performance in visual search tasks. i. contrast detection and discrimination. *Journal of Optical Society of America*, 20:1341 – 1355, 2003.

[5] N. J. Butko and J. R. Movellan. Infomax control of eye movements. *IEEE Trans. on Auton. Ment. Dev.*, 2(2):91–107, June 2010.

[6] N. Dalal and B. Triggs. Histograms of oriented gradients for human detection. In *Conf. on Computer Vision and Pattern Recognition*, pages 886–893, 2005.

[7] T. Dean, M. A. Ruzon, M. Segal, J. Shlens, S. Vijayanarasimhan, and J. Yagnik. Fast , Accurate Detection of 100,000 Object Classes on a Single Machine. In *Conf. on Computer Vision and Pattern Recognition*, 2013.

[8] S. Divvala, A. Efros, and M. Hebert. How important are "deformable parts" in the deformable parts model? In *European Conf. on Computer Vision*, Workshop on Parts and Attributes, pages 31–40, 2012.

[9] M. P. Eckstein. Visual search: a retrospective. *Journal of vision*, 11(5):14–, Jan. 2011.

[10] M. P. Eckstein, B. A. Drescher, and S. S. Shimozaki. Attentional cues in real scenes, saccadic targeting, and Bayesian priors. *Psychological science*, 17(11):973–80, Nov. 2006.

[11] J. Elder, S. Prince, Y. Hou, M. Sizintsev, and E. Olevskiy. Pre-Attentive and Attentive Detection of Humans in Wide-Field Scenes. *Int'l Journal of Computer Vision*, 72(1):4766, 2007.

[12] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The PASCAL Visual Object Classes Challenge 2007 (VOC2007) Results. http://www.pascal-network.org/challenges/VOC/voc2007/workshop/index.html.

[13] P. Felzenszwalb, R. Girshick, and D. McAllester. Cascade object detection with deformable part models. In *Conf. on Computer Vision and Pattern Recognition*, 2010.

[14] P. F. Felzenszwalb, R. B. Girshick, D. McAllester, and D. Ramanan. Object detection with discriminatively trained part based models. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 32(9):1627–1645, 2010.

[15] J. Freeman and E. P. Simoncelli. Metamers of the ventral stream. *Nature Neuroscience*, 14(9):1195–1201, 2011.

[16] T. Gao and D. Koller. Discriminative learning of relaxed hierarchy for large-scale visual recognition. In *Int'l Conf. on Computer Vision*, pages 2072–2079, 2011.

[17] R. B. Girshick, P. F. Felzenszwalb, and D. McAllester. Discriminatively trained deformable part models, release 5. http://people.cs.uchicago.edu/ rbg/latent-release5/.

[18] B. Hariharan, J. Malik, and D. Ramanan. Discriminative decorrelation for clustering and classification. In *European Conf. on Computer Vision*, 2012.

[19] H. Jegou, M. Douze, and C. Schmid. Product quantization for nearest neighbor search. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 33(1):117–128, Jan 2011.

[20] I. Kokkinos. Rapid deformable object detection using dual-tree branch-and-bound. In *Advances in Neural Information Processing*, 2011.

[21] I. Kokkinos. Bounding part scores for rapid detection with deformable part models. In *2nd Parts and Attributes Workshop, in conjunction with ECCV*, pages 41–50, 2012.

[22] P. Kontschieder, S. R. Bul, A. Criminisi, P. Kohli, M. Pelillo, and H. Bischof. Context-sensitive decision forests for object detection. In *Advances in Neural Information Processing*, 2012.

[23] C. H. Lampert. An efcient divide-and-conquer cascade for nonlinear object detection. In *Conf. on Computer Vision and Pattern Recognition*, 2010.

[24] C. H. Lampert, M. B. Blaschko, and T. Hofmann. Efficient subwindow search: A branch and bound framework for object localization. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 31(12):2129–2142, Dec 2009.

[25] H. Larochelle and G. Hinton. Learning to combine foveal glimpses with a third-order Boltzmann machine. In *Advances in Neural Information Processing*, pages 1–9, 2010.

[26] D. M. Levi, S. A. Klein, and A. P. Aitsebaomo. Vernier acuity, crowding and cortical magnification. *Vision Research*, 25(7):963–977, 1985.

[27] D. G. Lowe. Distinctive Image Features from Scale-Invariant Keypoints. *International Journal of Computer Vision*, 60(2):91–110, Nov. 2004.

[28] S. C. Mack and M. P. Eckstein. Object co-occurrence serves as a contextual cue to guide and facilitate visual search in a natural viewing environment. *Journal of vision*, 11(9):1–16, Jan. 2011.

[29] T. Malisiewicz, A. Gupta, and A. A. Efros. Ensemble of exemplar-svms for object detection and beyond. In *ICCV*, 2011.

[30] C. Morvan and L. T. Maloney. Human visual search does not maximize the post-saccadic probability of identifying targets. *PLoS computational biology*, 8(2):e1002342, Feb. 2012.

[31] J. Najemnik and W. S. Geisler. Optimal eye movement strategies in visual search. *Nature*, 434:387 – 391, 2005.

[32] J. Najemnik and W. S. Geisler. Simple summation rule for optimal fixation selection in visual search. *Vision research*, 49(10):1286–94, June 2009.

[33] M. B. Neider and G. J. Zelinsky. Scene context guides eye movements during visual search. *Vision research*, 46(5):614–21, Mar. 2006.

[34] M. Pedersoli, A. Vedaldi, and J. Gonzalez. A coarse-to-fine approach for fast deformable object detection. In *Conf. on Computer Vision and Pattern Recognition*, pages 1353–1360, 2011.

[35] T. J. Preston, F. Guo, K. Das, B. Giesbrecht, and M. P. Eckstein. Neural representations of contextual guidance in visual search of real-world scenes. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 33(18):7846–55, May 2013.

[36] N. Razavi, J. Gall, and L. Van Gool. Scalable multi-class object detection. In *Conf. on Computer Vision and Pattern Recognition*, pages 1505–1512, 2011.

[37] X. Ren and D. Ramanan. Histograms of sparse codes for object detection. In *Conf. on Computer Vision and Pattern Recognition*, 2013.

[38] L. W. Renninger, J. M. Coughlan, P. Verghese, and J. Malik. An information maximization model of eye movements. In *Advances in Neural Information Processing*, pages 1121–1128, 2004.

[39] J. Rovamo, L. Leinonen, P. Laurinen, and V. Virsu. Temporal integration and contrast sensitivity in foveal and peripheral vision. *Perception*, 13(6):665–74, Jan. 1984.

[40] M. Sadeghi and D. Forsyth. Fast Template Evaluation with Vector Quantization. In *Advances in Neural Information Processing*, 2013.

[41] T. Serre, L. Wolf, and T. Poggio. Object recognition with features inspired by visual cortex. In *Conf. on Computer Vision and Pattern Recognition*, 2005.

[42] H. O. Song, S. Zickler, T. Althoff, R. Girshick, M. Fritz, C. Geyer, P. Felzenszwalb, and T. Darrell. Sparselet Models for Efficient Multiclass Object Detection. In *European Conf. on Computer Vision*, 2012.

[43] H. Strasburger, I. Rentschler, and M. Jüttner. Peripheral vision and pattern recognition: a review. *Journal of vision*, 11(5):13, Jan. 2011.

[44] B. W. Tatler. The central fixation bias in scene viewing: selecting an optimal viewing position independently of motor biases and image feature distributions. *Journal of vision*, 7(14):4.1–17, Jan. 2007.

[45] A. Torralba, A. Oliva, M. S. Castelhano, and J. M. Henderson. Contextual guidance of eye movements and attention in real-world scenes: The role of global features in object search. *Psychological Review*, 113(4):766–786, 2006.

[46] K. E. A. van de Sande, J. R. R. Uijlings, T. Gevers, and A. W. M. Smeulders. Segmentation as selective search for object recognition. In *Int'l Conf. on Computer Vision*, 2011.

[47] P. Verghese. Active search for multiple targets is inefficient. *Vision Research*, 74:61–71, 2012.

[48] P. Viola and M. J. Jones. Robust real-time face detection. *Int'l Journal of Computer Vision*, 57(2):137–154, May 2004.

[49] T. Wertheim. Über die indirekte sehschärfe. *Zeitschrift für Psychologie und Physiologie der Sinnesorgane*, 7:172–187, 1894.

[50] G. J. Zelinsky. A theory of eye movements during target acquisition. *Psychological Review*, 115:787–835, 2008.

[51] S. Zhang and M. P. Eckstein. Evolution and optimality of similar neural mechanisms for perception and action during search. *PLoS Computational Biology*, 6(9):e1000930, 2010.

[52] W. Zhang, H. Yang, D. Samaras, and G. J. Zelinsky. A computational model of eye movements during object class detection. In *Advances in Neural Information Processing*, 2006.

[53] L. Zhu, Y. Chen, A. Yuille, and W. Freeman. Latent hierarchical structural learning for object detection. In *Conf. on Computer Vision and Pattern Recognition*, 2010.

[54] X. Zhu, C. Vondrick, D. Ramanan, and C. Fowlkes. Do we need more training data or better models for object detection? In *British Machine Vision Conf.*, pages 1–11, 2012.